

IN THE CLAIMS

Please amend the claims as follows.

1. (Currently Amended) A computing system comprising a plurality of nodes connected by a network, wherein the plurality of nodes include a group membership service operable to determine membership in a group formed by the plurality of nodes of a process executing on a node in the plurality of nodes for an application distributed across two or more of the plurality of nodes, said membership communicated between the plurality of nodes in the network utilizing a proposal message including data defining a plurality of relationships between the plurality of nodes and sent by a coordinator node for receipt by each node in the plurality of nodes and a commit message sent by the coordinator node to each node in the plurality of nodes after receiving acknowledgement that the proposal message has reached each node of the plurality of nodes, and further wherein the plurality of nodes communicate with each other to detect a failure on a first node of the plurality of nodes and to transfer applications from the first node to other nodes of the plurality of nodes in the group on detecting the failure.

2. (Currently Amended) A method of maintaining high availability in a server cluster having a plurality of nodes, the method comprising:

instantiating a group communications service, a group membership service and a system resource manager on each node of the plurality of nodes, the plurality of nodes forming a group;

communicating process membership in the group utilizing a proposal message sent by a coordinator node for receipt by each node in the plurality of nodes and a commit message sent by the coordinator node to each node in the plurality of nodes after receiving acknowledgement that the proposal message has reached each node of the plurality of nodes, wherein the proposal message includes data defining one or more relationships between the plurality of nodes;

communicating between the group communications service, the group membership service and the system resource manager on each node of the plurality of nodes to detect process failures and node failures within the group;

upon detecting a failure in a first node of the plurality of nodes, transferring applications to other nodes of the plurality of nodes; and

updating, by the group membership service, process membership in a distributed application upon detecting a process failure on a node of the plurality of nodes.

3. (Currently Amended) A computer-readable medium having instructions thereon, wherein the instructions, when executed in a computer, perform a method comprising:

instantiating a group communications service, a group membership service and a system resource manager on each node of a plurality of nodes, the plurality of nodes forming a group;

communicating process membership in a group utilizing a proposal message including data defining one or more relationships between the plurality of nodes sent by a coordinator node for receipt by each node in the plurality of nodes and a commit message sent by the coordinator node to each node in the plurality of nodes after receiving acknowledgement that the proposal message has reached each node of the plurality of nodes;

communicating between the group communications service, the group membership service and the system resource manager on each node of the plurality of nodes to detect process failures and node failures within the group;

upon detecting a failure in a first node of the plurality of nodes, transferring applications to other nodes of the plurality of nodes; and

updating, by the group membership service, process membership in a distributed application upon detecting a process failure on a node of the plurality of nodes.

4. (Previously Presented) The computing system of claim 1, wherein the plurality of nodes includes an initiator node to send the proposal message to the coordinator node.

5. (Previously Presented) The computing system of claim 4, wherein the coordinator node comprises a longest running node in the plurality of nodes.

6. (Previously Presented) The computing system of claim 4, wherein the plurality of nodes are arranged in a network ring, the order of the plurality of nodes in the network ring being

defined by a cluster membership age of each node in the plurality of nodes and wherein the coordinator node forwards the proposal message to a first node of the plurality of nodes,

and wherein the proposal message is forwarded by a receiving node in the network ring to a successor node of the receiving node.

7. (Previously Presented) The computing system of claim 6, wherein the coordinator node issues the commit message upon receiving the proposal message from a non-initiator node in the network ring.

8. (Previously Presented) The method of claim 2, wherein communicating the proposal message includes sending by an initiator node the proposal message to the coordinator node.

9. (Previously Presented) The method of claim 8, wherein the coordinator node comprises a longest running node in the plurality of nodes.

10. (Previously Presented) The method of claim 8, further comprising:
arranging the plurality of nodes in a network ring;
forwarding by the coordinator node the proposal message to a first node of the plurality of nodes; and
forwarding by the first node to a next node in the network ring.

11. (Previously Presented) The method of claim 10, wherein the coordinator node issues the commit message upon receiving the proposal message from a non-initiator node in the ring.

12. (Previously Presented) The method of claim 10, wherein upon receiving the commit message a node of the plurality of nodes in the network ring performs the tasks of:
caching the commit message;
forwarding the commit message to a next node in the network ring;

upon completing forwarding the commit message delivering the commit message to each process of a process group on the node.

13. (Previously Presented) The computer readable medium of claim 3, wherein communicating the proposal message includes sending by an initiator node the proposal message to the coordinator node.

14. (Previously Presented) The computer readable medium of claim 13, wherein the coordinator node comprises a longest running node in the plurality of nodes.

15. (Previously Presented) The computer readable medium of claim 13, wherein the method further comprises:

arranging the plurality of nodes in a network ring;
forwarding by the coordinator node the proposal message to a first node of the plurality of nodes; and
forwarding by the first node to a next node in the network ring.

16. (Previously Presented) The computer readable medium of claim 15, wherein the coordinator node issues the commit message upon receiving the proposal message from a non-initiator node in the ring.

17. (Previously Presented) The computer readable medium of claim 15, wherein upon receiving the commit message a node of the plurality of nodes in the network ring performs the tasks of:

 caching the commit message;
 forwarding the commit message to a next node in the network ring;
 upon forwarding the commit message delivering the commit message to each process of a process group on the node.